Machine Learning manuscript No.
(will be inserted by the editor)

Asymptotic consistency and order specification for logistic classifier chains in multi-label learning

Paweł Teisseyre

Received: date / Accepted: date

arXiv:1602.07466v1 [cs.LG] 24 Feb 2016

Abstract Classifier chains are popular and effective method to tackle a multi-label classification problem. The aim of this paper is to study the asymptotic properties of the chain model in which the conditional probabilities are of the logistic form. In particular we find conditions on the number of labels and the distribution of feature vector under which the estimated mode of the joint distribution of labels converges to the true mode. Best of our knowledge, this important issue has not yet been studied in the context of multi-label learning. We also investigate how the order of model building in a chain influences the estimation of the joint distribution of labels. We establish the link between the problem of incorrect ordering in the chain and incorrect model specification. We propose a procedure of determining the optimal ordering of labels in the chain, which is based on using measures of correct specification and allows to find the ordering such that the consecutive logistic models are best possibly specified. The other important question raised in this paper is how accurately can we estimate the joint posterior probability when the ordering of labels is wrong or the logistic models in the chain are incorrectly specified. The numerical experiments illustrate the theoretical results.

Keywords classifier chains · logistic regression · joint mode estimation · label ordering · asymptotic consistency

1 Introduction

In multi-label classification the task is to automatically assign an object to multiple categories based on its characteristics. Each object of our interest is described by a feature vector $\mathbf{x}$ belonging to p-dimensional space and vector of K labels $\mathbf{y} = (y_1, \ldots, y_K)'$. In this paper we consider binary labels such that $y_k = 1$ indicates that the considered object belongs to k-th category or has the k-th property. The issue has recently attracted significant attention, motivated by an increasing number of applications such as image and video annotation

Paweł Teisseyre
Institute of Computer Science, Polish Academy of Sciences
Jana Kazimierza 5 01-248 Warsaw, Poland
Tel.: +48-22-380-05-55
E-mail: teisseyrep@ipipan.waw.pl

(assigning metadata to digital images or videos), music categorization (assigning various emotions to the songs), text categorization (assigning various subjects to documents), direct marketing (predicting which products will be purchased), functional genomics (determining the functions of genes and proteins) and medical diagnosis (predicting types of diseases).

The key problem in multi-label learning is how to utilize label dependencies to improve the classification performance, motivated by which number of multi-label algorithms have been proposed in recent years (see Madjarov et al. (2012) for extensive comparison of several methods). A naive approach, called binary relevance (BR), which is based on building separate classification models for each label and which does not take into account any dependencies between labels, is deficient in many applications. A natural approach, to tackle the issue, is to model the joint conditional probability $P(y_1,\ldots,y_n|\mathbf{x})$ and then to choose the most probable set of labels for a given $\mathbf{x}$, i.e. the mode of the joint distribution $\mathbf{y}^*(\mathbf{x}) = \arg\max_{\mathbf{y}\in\{0,1\}^K} p(\mathbf{y}|\mathbf{x})$. This approach minimizes the subset 0/1 loss (Dembczyński et al. (2012)), which generalizes the well-known 0/1 loss from the conventional to the multi-label setting. Direct modelling of the joint distribution is usually problematic. Classifier chain method (Read et al. (2011), Dembczyński et al. (2010)) fixes this problem by using chain probability rule

$$P(\mathbf{y}|\mathbf{x}) = P(y_1,\ldots,y_n|\mathbf{x}) = \prod_{k=1}^{K} P(y_k|y_1,\ldots,y_{k-1},\mathbf{x}), \tag{1}$$

which allows to estimate the conditional probabilities in (1) by using single label classifiers in which y_k is treated as response variable, whereas $y_1,\ldots,y_{k-1},\mathbf{x}$ are treated as explanatory variables. This gives the approximations of $P(y_k|y_1,\ldots,y_{k-1},\mathbf{x})$ and thus also the approximation of the joint distribution, which will be denoted by $\hat{P}(\mathbf{y}|\mathbf{x})$. Classifier chains are commonly assigned to the group of problem-transformation methods (Tsoumakas and Katakis (2007)), as in this case the multi-label problem is transformed into several single label problems. One of the main difficulties of classifier chains is their sensitivity to the ordering of labels during building the models. Although the order of conditioning in (1) is not relevant, the order of training models can affect significantly the estimation accuracy.

In this paper we consider the classifier chains combined with logistic model fitting, i.e. logistic models are used to estimate the conditional probabilities in (1). Our motivation for studying this particular combination is that the logistic model is one of the most commonly used classifiers in practice. The combination of classifier chains and logistic regression has already been used by many authors, in various applications, e.g. in Kumar et al. (2013), Dembczyński et al. (2010), Dembczyński et al. (2012), Montañes et al. (2014). However, there are no theoretical results on large sample properties for such procedures. This paper intends to fill this gap. Moreover, the logistic model has strong theoretical background (Fahrmeir and Kaufmann (1985), Fahrmeir (1987)), which allows to prove some asymptotic results for classifier chains as well.

First, we study the asymptotic properties of classifier chains. We impose the conditions for the distribution of the feature vector $\mathbf{x}$ and number of labels K under which the estimated mode $\hat{\mathbf{y}}(\mathbf{x}) = \arg\max_{\mathbf{y}\in\{0,1\}} \hat{P}(\mathbf{y}|\mathbf{x})$ converges in probability to the true mode $\mathbf{y}^*(\mathbf{x})$. The above property can be proved, provided that there exists the "optimal" ordering of labels, such that the consecutive conditional probabilities in (1) are of the logistic form. However, in practice, the "optimal" ordering is usually unknown or it may not exist at all (this is the case when the data generation scheme is different from logistic one). This leads to further important questions discussed in this paper. What is the influence of the ordering of labels on the estimation of the joint distribution (and thus on the estimation of the mode)? How

accurate can we estimate the conditional probabilities in (1) when the ordering of labels is wrong or the logistic model is incorrectly specified? It turns out that in this case we attempt to estimate the parameters of assumed logistic model which is closest to the true one in the sense of average Kullback-Leibler distance. Moreover, we propose a procedure of finding the optimal ordering of labels, which uses measures of correct specification. The procedure allows to determine the ordering such that the consecutive conditional probabilities in (1) are best possibly specified. Although the following results are obtained for logistic model, we believe that some ideas can be expanded to the case when other base-classifiers are used. The problem of label ordering has been discussed in literature, however the direct influence on the joint mode estimation has not been investigated. Kumar et al. (2013) proposed to use log-likelihood function as a measure of correct ordering and beam search to reduce the large number of possible orderings. To eliminate an effect of ordering of labels, Read et al. (2009) proposed to average the multi-label predictions over randomly chosen set of permutations. This extension, called ensembled classifier chain (ECC) improves the results of single chain, however it does not answer the question how to find the optimal ordering.

The rest of the paper is organized as follows. In Section 2 we describe the logistic classifier chain model (LCC), which is motivated by (1). In Section 3 we state the Theorem concerning the consistency of the joint mode estimation. The problem of ordering the labels is discussed in Section 4. In Section 4 we also review the measures of correct specification and the procedure of determining the optimal ordering is presented. In Section 5, problem of inference in classifier chain model is described. Finally, the results of selected numerical experiments are reported in Section 6. Section 7 contains final conclusions and 8 contains proofs.

2 Logistic classifier chain (LCC) model

In multi-label learning each object of our interest is described by a pair $(\mathbf{x},\mathbf{y})$, where $\mathbf{x}\in R^p$ is the random vector of p explanatory variables (features) and $\mathbf{y}=(y_1,\ldots,y_{K_n})'$ is the vector of K_n binary responses (labels). The first coordinate of $\mathbf{x}$ is equal 1, which corresponds to the intercept. We assume that the number of responses depends on the number of observations, which reflects the common situation of large number of labels. Moreover, this assumption can correspond to the situation of sparse labels (equal zero for the majority of cases) - we add the new label at a time when the first non-zero value of this label appears. We assume the following data generation scheme. First, we assume that there exists a permutation/ordering of labels $\pi^*(1),\ldots,\pi^*(K_n)$ such that the probability of $y_{\pi^*(k)}=1$, given $\mathbf{x}$ and the previous labels $y_{\pi^*(1)},\ldots,y_{\pi^*(k-1)}$, is always of the logistic form, i.e.

$$P(y_{\pi^*(k)}=1|\mathbf{x},y_{\pi^*(1)},\ldots,y_{\pi^*(k-1)})=\sigma(\mathbf{z}_k(\pi^*)'\theta_k(\pi^*)),\quad \text{for all } k=1,\ldots,K_n, \tag{2}$$

where $\sigma(s)=1/(1+\exp(-s))$ is a logistic function; $\mathbf{z}_k(\pi^*)=(\mathbf{x}',y_{\pi^*(1)},\ldots,y_{\pi^*(k-1)})'$ is a combined vector of features and labels with indices $\pi^*(1),\ldots,\pi^*(k-1)$; $\theta_k(\pi^*)\in R^{p+k-1}$ is a corresponding vector of parameters. In total we have $pK_n+K_n(K_n-1)/2$ unknown parameters. For notation simplicity we will assume that the ordering of labels, for which (2) holds, is an identity permutation $\pi^*(k)=k$ and in this case we will write in brief $\mathbf{z}_k$ and θ_k instead of $\mathbf{z}_k(\pi^*)$ and $\theta_k(\pi^*)$. So, in this case, we can simply write

$$P(y_k=1|\mathbf{x},y_1,\ldots,y_{k-1})=\sigma(\mathbf{z}_k'\theta_k),\quad \text{for all } k=1,\ldots,K_n. \tag{3}$$

The joint distribution is of the form

$$P(\mathbf{y}|\mathbf{x}) = P(y_1,\ldots,y_k|\mathbf{x}) = \prod_{k=1}^{K_n} P(y_k = 1|\mathbf{x}, y_1,\ldots,y_{k-1}) = \prod_{k=1}^{K_n} \sigma(\mathbf{z}_k'\theta_k)^{y_k}[1-\sigma(\mathbf{z}_k'\theta_k)]^{1-y_k}. \tag{4}$$

The joint distribution corresponding to BR model, with the conditional distributions of the logistic form, can be written in an analogous form to (4), with the only difference that $\mathbf{z}_k = \mathbf{x}$. Diagrams on Figure 1 represent the networks corresponding to LCC model and BR model.

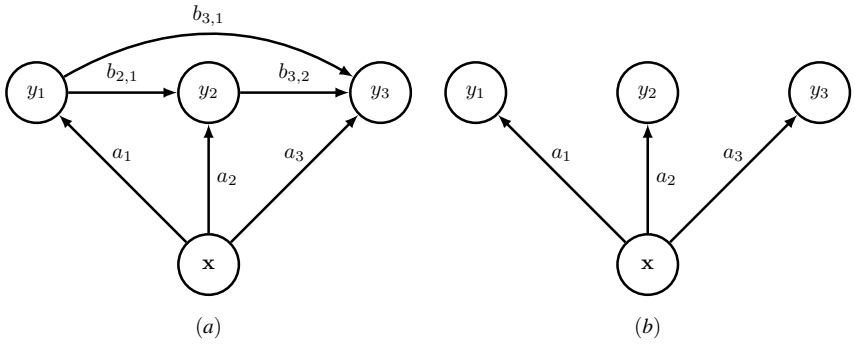

Fig. 1 Diagrams representing data generation schemes in LCC model (a) and BR model (b), in the case of $K_n = 3$ labels. The following notation is used: $\theta_k = (a_k', b_k')'$, where $a_k \in R^p$ is a subvector of parameters corresponding to $\mathbf{x}$ and $b_k = (b_{k,1},\ldots,b_{k,k-1})' \in R^{k-1}$ is a subvector of parameters corresponding to $k-1$ first labels in the chain.

As the true joint distribution is usually unknown, it is estimated based on training examples $\mathscr{D} = (\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$, $i = 1,\ldots,n$ ($\mathbf{x}^{(i)} \in R^p$, $\mathbf{y}^{(i)} \in \{0,1\}^{K_n}$) which are assumed to be generated from (3). Observe that conditional distribution (3) can be estimated by fitting logistic regression model (on the set $\mathscr{D}$) in which y_k is a response variable and $\mathbf{z}_k = (\mathbf{x}', \mathbf{y}_{1:(k-1)}')'$ is a vector of explanatory variables. By $\mathscr{M}(y_k, \mathbf{z}_k)$ we will denote a logistic model fitted on set $\mathscr{D}$, in which y_k is a response variable and $\mathbf{z}_k$ is a vector of explanatory variables.

Let $l_{y_k,\mathbf{z}_k}(\cdot)$ be the log-likelihood function calculated for model $\mathscr{M}(y_k,\mathbf{z}_k)$ using $\mathscr{D}$ and $\hat{\theta}_k = \arg\max_\theta l_{y_k,\mathbf{z}_k}(\theta)$ be the maximum likelihood estimator of θ_k. By building K_n independent logistic models $\mathscr{M}(y_1,\mathbf{z}_1),\ldots,\mathscr{M}(y_{K_n},\mathbf{z}_{K_n})$, one can estimate parameters $\theta_1,\ldots,\theta_{K_n}$ in the consecutive models in the chain and thus also the joint distribution

$$\hat{P}(\mathbf{y}|\mathbf{x}) = \prod_{k=1}^{K_n} \sigma(\mathbf{z}_k'\hat{\theta}_k)^{y_k}[1-\sigma(\mathbf{z}_k'\hat{\theta}_k)]^{1-y_k}. \tag{5}$$

3 Consistency of the joint mode estimation

In the following we will assume that the true ordering of labels $(\pi^*(1),\ldots,\pi^*(K_n)) = (1,\ldots,K_n)$, for which (2) holds, is known and the logistic models $\mathscr{M}(y_1,\mathbf{z}_1),\ldots,\mathscr{M}(y_{K_n},\mathbf{z}_{K_n})$ are fitted according to this order based on training set $\mathscr{D}$.

Denote by $\mathbf{y}^*(\mathbf{x}) = \arg\max_{\mathbf{y}\in\{0,1\}^{K_n}} P(\mathbf{y}|\mathbf{x})$ the true joint mode of the distribution (4), for some new observation $\mathbf{x} \in R^p$ and by $\hat{\mathbf{y}}(\mathbf{x}) = \arg\max_{\mathbf{y}\in\{0,1\}^{K_n}} \hat{P}(\mathbf{y}|\mathbf{x})$ the estimated mode for

point $\mathbf{x}$, calculated based on set $\mathscr{D}$. As in multi-label problems we are usually interested in finding the true joint mode, it is worthwhile to approximate it accurately. In this section we will give the assumptions under which the true joint mode can be estimated consistently.

Let $s_{y_k,\mathbf{z}_k}(\theta) = \frac{\partial}{\partial\theta} l_{y_k,\mathbf{z}_k}(\theta)$ be the gradient and $H_{y_k,\mathbf{z}_k}(\theta) = -\frac{\partial^2}{\partial\theta\partial\theta'} l_{y_k,\mathbf{z}_k}(\theta)$ be the negative Hessian matrix, based on model $\mathscr{M}(y_k,\mathbf{z}_k)$ and data $\mathscr{D}$. Let Z_k be the matrix whose i-th row is $\mathbf{z}_k^{(i)} = (\mathbf{x}^{(i)'}, \mathbf{y}_{1:(k-1)}^{(i)'})'$, for $i = 1,\ldots,n$. For logistic model we have

$$s_{y_k,\mathbf{z}_k}(\theta) = Z_k'(\mathbf{w}_k - E(\mathbf{w}_k|Z_k)),$$

where $\mathbf{w}_k = (y_k^{(1)},\ldots,y_k^{(n)})'$ is a response vector corresponding to k-th label. To prove our main Theorem, we use the following Lemma, which is an important technical tool. The Lemma follows from Zhang (2009), after noting that binary random variable w satisfies $E\exp(t(w-Ew)) \leq \exp(t^2/8)$ and taking $\sigma = 1/2$, $\eta = (\varepsilon+1/2)^2$ in his Proposition 10.2.

Lemma 1 *Let* $\mathbf{w}$ *be n-dimensional vector consisting of independent binary variables having not necessarily the same distribution and* Z *be* $n \times q$ *fixed matrix. Then for any* $\eta > 1$ *we have*

$$P[||Z'(\mathbf{w}-E\mathbf{w})||^2 \geq tr(Z'Z)\eta] \leq \exp(-\eta/20).$$

The inequality in Lemma 1 holds also for random Z, which can be easily seen by conditioning

$$P[||Z'(\mathbf{w}-E\mathbf{w})||^2 \geq tr(Z'Z)\eta] = E\{P[||Z'(\mathbf{w}-E(\mathbf{w}|Z))||^2 \geq tr(Z'Z)\eta|Z]\} \leq \exp(-\eta/20),$$

where the first expectation in above formula is taken with respect to Z. It is seen that Lemma 1 can be used to bound from above $||s_{y_k,\mathbf{z}_k}(\theta)||$, which is an important part of the proof of Theorem 1. In the following Theorem, we impose conditions under which the true mode $\mathbf{y}^*(\mathbf{x})$ is estimated consistently. For the proof, see in Section 8.1.

Theorem 1 *Let* $\mathbf{x}$ *be a fixed point at which the joint mode is calculated. Assume that:*

1. *there exists* $\varepsilon > 0$*, such that* $P[\mathbf{y}^*(\mathbf{x})|\mathbf{x}] > P[\mathbf{y}|\mathbf{x}] + \varepsilon$*, for all* $\mathbf{y} \neq \mathbf{y}^*(\mathbf{x})$*,*
2. $E||\mathbf{x}^{(i)}||^2 < \infty$*,*
3. $K_n^4 \log(K_n)/n \to 0$ *and* K_n *is monotonic function of n,*
4. *there exists constant* c_1 *such that*

$$P\left[\min_{1\leq k\leq K_n} \lambda_{\min}[H_{y_k,\mathbf{z}_k}(\theta_k)/n] \geq c_1\right] \to 1,$$

where $\lambda(\cdot)$ *is an eigenvalue (*$\lambda_{\min}$ *is minimal eigenvalue).*

Then, we have that $\hat{\mathbf{y}}(\mathbf{x})$ *is consistent, i.e.* $P[\hat{\mathbf{y}}(\mathbf{x}) = \mathbf{y}^*(\mathbf{x})] \to 1$*, as* $n \to \infty$*.*

The first assumption ensures that the true mode $\mathbf{y}^*(\mathbf{x})$ is unique for $\mathbf{x}$. The second assumption on the distribution of $\mathbf{x}^{(i)}$ is rather mild and is satisfied for the majority of common distributions. The third one indicates that the number of possible labels should be small in comparison with the number of training examples n. It can be verified that if for almost every $\mathbf{y}$, $|\{k : y_k \neq 1\}| \leq C < \infty$, with probability tending to 1, then the third assumption can be weakened to $K_n^3 \log(K_n)/n \to 0$. This would reflect the situation in which the number of possible labels is large, however only a few of them can be equal to 1 simultaneously. The last assumption is a regularity condition, analogous to assumption (A4) used in Chen and Chen (2012) to prove the consistency of Extended Bayesian Information Criterion (EBIC)

for logistic regression model. Conditions similar to our last assumption are used in many papers devoted to asymptotic results in logistic regression (see e.g. in Qian and Field (2002) and Mielniczuk and Szymanowski (2015)). Observe that,

$$\lambda_{\min}[H_{y_k,\mathbf{z}_k}(\theta_k)/n] \geq \min_i[\sigma(\mathbf{z}_k^{(i)'}\theta_k)(1-\sigma(\mathbf{z}_k^{(i)'}\theta_k))]\lambda_{\min}(Z_k'Z_k/n)$$

and thus condition 4 is implied by the following two assumptions:

$$P\left[\min_{1\leq k\leq K_n}\min_i[\sigma(\mathbf{z}_k^{(i)'}\theta_k)(1-\sigma(\mathbf{z}_k^{(i)'}\theta_k))] \geq c_{11}\right] \to 1, \tag{6}$$

and

$$P\left[\min_{1\leq k\leq K_n}\lambda_{\min}[Z_k'Z_k/n] \geq c_{12}\right] \to 1, \tag{7}$$

for some constants c_{11} and c_{12}. Above condition (6) indicates that conditional variance of y_k given $\mathbf{x}, y_1, \ldots, y_{k-1}$ must be bounded from 0, for all observations and all tasks. The convergence in (7) is regularity condition on the design matrices. Note that (7) is weaker than assumption stating that $Z_k'Z_k/n$ converges to some positive definite matrix, which is commonly used assumption in regression analysis (see e.g. Nishii (1984)).

Using the proof of Theorem 1, we can also show the consistency of parameter estimators. Let $\theta = (\theta_1, \ldots, \theta_{K_n})' \in R^{pK_n+K_n(K_n-1)/2}$ be the vector of parameters corresponding to all models in the chain and let $\hat{\theta} = (\hat{\theta}_1, \ldots, \hat{\theta}_{K_n})' \in R^{pK_n+K_n(K_n-1)/2}$ be the corresponding vector of estimators. The following Corollary holds (see Section 8.2 for the proof).

Corollary 1 *Assume that conditions 2,3,4 from Theorem 1 hold. Then for any* $\varepsilon > 0$

$$P(||\hat{\theta} - \theta|| > \varepsilon K_n^{-1}) \to 0,$$

as $n \to \infty$.

4 Ordering of labels in classifier chain

In practice, the true ordering of labels $(\pi^*(1), \ldots, \pi^*(K_n)) = (1, \ldots, K_n)$, for which (2) holds, is unknown. Moreover, it may happen that the ordering, for which the consecutive conditional probabilities are of the logistic form, does not exist at all. In this section we study the influence of the ordering of labels on the estimation of the joint distribution (4).

To illustrate the problem, consider the following example on real dataset `emotions` (Trohidis et al. (2008)), having 6 binary labels and 72 features. To obtain convenient visualization, we initially replaced the features by their first principal component, which explains about 88% of variance of the original features. Figure 4 shows estimated joint distributions (for 4 selected labellings) as functions of the first principal component. To estimate joint distributions we use classifier chains with logistic regression and two orders of fitting models: $(1,2,3,4,5,6)$ and $(6,5,4,3,2,1)$. Obviously in this case the true ordering of labels is unknown, however it is clearly seen that the order of fitting models in the chain affects the estimate of the joint probability significantly.

Let π^0 be assumed ordering of labels and recall that $\mathbf{z}_k(\pi^0) = (\mathbf{x}', y_{\pi^0(1)}, \ldots, y_{\pi^0(k-1)})'$. We consider the situation when logistic models $\mathscr{M}(y_{\pi^0(k)}, \mathbf{z}_k(\pi^0))$, for $k = 1, \ldots, K_n$ are fitted according to π^0, whereas the data is generated according to the true ordering π^*. Unfortunately, in this case we may encounter the problem of incorrect model specification, e.g.

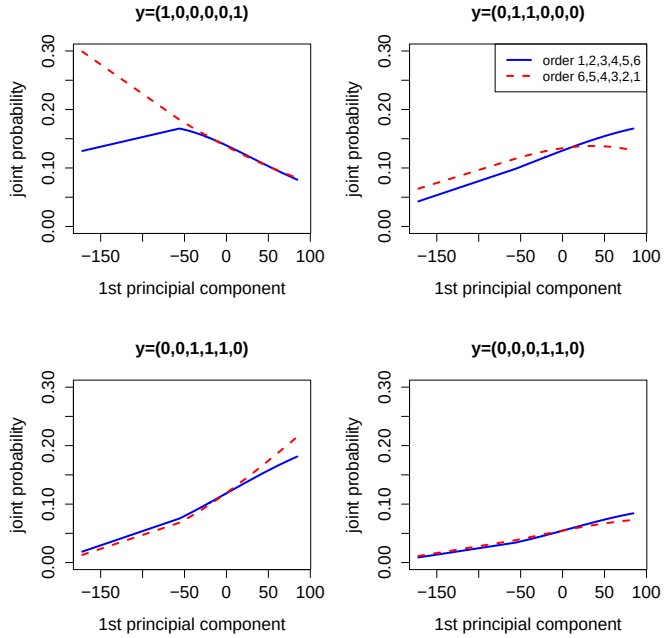

Fig. 2 Estimated joint distributions (for 4 combinations of labels) as functions of the 1st principal component, for `emotions` dataset.

we fit the logistic model whereas the true function, from which the data is generated, may not be necessarily a logistic one. To assess the problem, consider the simple case of 2 labels and one feature $x \in R$. Assume that the true ordering is $\pi^* = (1,2)$ and the labels are generated as $P(y_1 = 1|x) = \sigma(x)$ and $P(y_2 = 1|x,y_1) = \sigma(x + ay_1)$, where $a \in R$ is a parameter. The joint distribution in this case is of the form $P(y_1,y_2|x) = \sigma(x)^{y_1}(1-\sigma(x))^{1-y_1}\sigma(x + ay_1)^{y_2}(1-\sigma(x+ay_1))^{1-y_2}$. Assume now that we fit logistic models according to reverse ordering $\pi^0 = (2,1)$, so we use logistic model to approximate $P(y_2 = 1|x)$ which is

$$P(y_2 = 1|x) = P(y_1 = 0, y_2 = 1|x) + P(y_1 = 1, y_2 = 1|x) = \sigma(x) + \sigma(x)[\sigma(x+a) - \sigma(x)].$$

It is easy to verify that there do not exist parameters $v,w \in R$ such that for all x, $P(y_2 = 1|x) = \sigma(v + wx)$, so we cannot approximate accurately the true distribution $P(y_2 = 1|x)$ using logistic model.

Recall that $\hat{\theta}_k(\pi^0)$ is a maximum likelihood estimator from fitted model $\mathscr{M}(y_{\pi^0(k)}, \mathbf{z}_k(\pi^0))$. We show what is the limit of $\hat{\theta}_k(\pi^0)$ in the general situation, when the assumed ordering π^0 is not necessarily the true one. Hjort and Pollard (1993) studied the performance of logistic model under incorrect model specification. See also Czado and Santner (1992) and Hjort (1988), where various implications for statistical inference are discussed. Using their methodology (see Section 5B in Hjort and Pollard (1993)) it can be shown that for any $\varepsilon > 0$

$$P(|\hat{\theta}_k(\pi^0) - \tilde{\theta}_k(\pi^0)| > \varepsilon) \to 0, \tag{8}$$

as $n \to \infty$, where

$$\tilde{\theta}_k(\pi^0) := \arg \min_{\gamma \in R^{p+k-1}} E\left\{ KL\left[P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0)), \sigma(\mathbf{z}_k(\pi^0)'\gamma) \right] \right\},$$

and $KL(q_1,q_2) := q_1 \log[q_1/q_2] + (1-q_1)\log[(1-q_1)/(1-q_2)]$ is a Kullback-Leibler divergence. The expectation in (8) is taken with respect to random vector $\mathbf{z}_k(\pi^0)$. So the limit $\tilde{\theta}_k(\pi^0)$ minimizes the expected KL divergence from the true distribution $P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0))$ to the postulated one with parameter $\sigma(\mathbf{z}_k(\pi^0)'\tilde{\theta}_k(\pi^0))$. The convergence in (8) can be interpreted as follows: we find the most accurate approximation of the true probability $P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0))$ with respect to the KL divergence, within the space of the logistic models. When $\pi^0 = \pi^*$, the true probability $P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0))$ is of the logistic form and then $\tilde{\theta}_k(\pi^0) = \theta_k(\pi^0)$. Moreover, we state the proposition which indicates that the expectations of the true distribution and the postulated one, with parameter $\tilde{\theta}_k(\pi^0)$, coincide.

Proposition 1 *The following equalities hold*

$$E[P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0))\mathbf{z}_k(\pi^0)] = E[\sigma(\mathbf{z}_k(\pi^0)'\tilde{\theta}_k(\pi^0))\mathbf{z}_k(\pi^0)] \tag{9}$$

and

$$E[P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0))] = E[\sigma(\mathbf{z}_k(\pi^0)'\tilde{\theta}_k(\pi^0))]. \tag{10}$$

Proof Using the fact that $\partial\sigma(\mathbf{z}_k(\pi^0)'\gamma)/\partial\gamma = \sigma(\mathbf{z}_k(\pi^0)'\gamma)(1-\sigma(\mathbf{z}_k(\pi^0)'\gamma))$, it is easy to see that

$$E\left[\frac{\partial KL\left[P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0)), \sigma(\mathbf{z}_k(\pi^0)'\gamma)\right]}{\partial\gamma}\right] =$$
$$E[-P(y_{\pi^0(k)} = 1|\mathbf{z}_k(\pi^0))\mathbf{z}_k(\pi^0) + \sigma(\mathbf{z}_k(\pi^0)'\gamma)\mathbf{z}_k(\pi^0)],$$

which yields (9). Equality (10) follows from (9) after noting that the first coordinate of $\mathbf{z}_k(\pi^0)$, corresponding to intercept, is equal 1. □

Hereinafter, we propose a method which aims to find the ordering such that the consecutive logistic models in the chain are correctly specified. We propose to use a forward procedure such that at each stage we choose the label corresponding to best specified logistic model. Note that we are more interested in evaluation of goodness of specification than evaluation of the goodness of fit. In order to have such a procedure, it is necessary to define a measure which indicates whether the given model is correctly specified or not.

4.1 Goodness of specification measures

For convenience, in this section we will focus on the link functions, which relate posterior probabilities to the linear combinations of the features. For the logistic model, the link function is the inverse of $\sigma(\cdot)$. We discuss the situation of fitting (based on response variable y and some input features $\mathbf{z}$) a logistic model having link function $g_0(\mu) = \log(\frac{\mu}{1-\mu})$ when in fact, the correct link function is $g_*(\mu)$. The idea is to define more general family of link functions $\mathscr{L} := \{g(\mu;\alpha) : \alpha \in R\}$, depending on some parameter α (or more parameters), in such a way that both assumed link and correct link are the members of $\mathscr{L}$. Having such a family, we may write $g_0(\mu) = g(\mu;\alpha_0)$ and $g_*(\mu) = g(\mu;\alpha_*)$, where α_0 is known whereas α_* is unknown. Using the first-order Taylor expansion about the assumed link we have the approximation

$$g_*(\mu) \approx g_0(\mu) + (\alpha_* - \alpha_0)\frac{\partial}{\partial\alpha}g(\mu;\alpha)|_{\alpha=\alpha_0}.$$

Table 1 Carriers for computing deviance $D(y,\mathbf{z})$.

Method	Carriers
Preigbon	$w_1 = 0.5(\log^2(\hat{\mu}/n) - \log^2(1-\hat{\mu}/n))$; $w_2 = -0.5(\log^2(\hat{\mu}/n) + \log^2(1-\hat{\mu}/n))$
Stukel	$w_1 = 0.5(\mathbf{z}'\hat{\theta})^2\mathbf{I}(\mathbf{z}'\hat{\theta} \geq 0)$; $w_2 = -0.5(\mathbf{z}'\hat{\theta})^2\mathbf{I}(\mathbf{z}'\hat{\theta} < 0)$
Prentice	$w_1 = -(1-\hat{\mu})^{-1}\log(\hat{\mu})$; $w_2 = -\hat{\mu}^{-1}\log(1-\hat{\mu})$
Guerrero-Johnson	$w_1 = 0.5(\mathbf{z}'\hat{\theta})^2$
Morgan	$w_1 = (\mathbf{z}'\hat{\theta})^3$
Aranda (asymmetric)	$w_1 = 1 + \hat{\mu}^{-1}\log(1-\hat{\mu})$

Observe that $g_*(\mu) = \mathbf{z}'\theta$, where $\mathbf{z}$ is a vector of input variables and θ is vector of unknown parameters associated with the unknown link and thus the assumed link can be approximated by

$$g_0(\mu) \approx \mathbf{z}'\theta + w(\alpha_0 - \alpha_*), \tag{11}$$

where $w := \frac{\partial}{\partial\alpha} g(\mu;\alpha)|_{\alpha=\alpha_0}$ is so-called carrier variable and $\gamma := (\alpha_0 - \alpha_*)$ is an unknown parameter. Variable w is unknown but we can approximate it in the following way. We initially fit model with logistic link $g_0(\mu)$ and estimate parameters using maximum likelihood method. This yields maximum likelihood estimates $\hat{\theta}$ and thus also $\hat{\mu}$, from which we can approximate $\hat{w} := \frac{\partial}{\partial\alpha} g(\hat{\mu};\alpha)|_{\alpha=\alpha_0}$. Now as a measure of correct link specification, we can use deviance statistic, defined as

$$D_{\hat{w}}(y,\mathbf{z}) := 2\{l_{y,(\mathbf{z},\hat{w})}(\hat{\theta}_e, \hat{\gamma}_e) - l_{y,\mathbf{z}}(\hat{\theta})\}, \tag{12}$$

where $\hat{\theta}_e, \hat{\gamma}_e$ denote estimators based on extended model with an additional variable $\hat{w}$. This requires merely refitting the original logistic model with additional variable $\hat{w}$. Observe that $\gamma = 0$ if the link is correctly specified. Small value of the deviance indicates that the link is correctly specified, whereas significant departure from 0 suggest incorrect specification. The above reasoning can be easily generalized to the case of more than 1 carrier, it simply requires refitting the original logistic model with additional $2,3,\ldots$ carriers. The crucial in the above idea is the choice of the family $\mathscr{L}$. Usually the proposed families are parametrized by one or two parameters. In the literature, various approaches have been explored. Best of our knowledge, the first attempt was made by Preigbon (1980) who use the family defined by

$$\mathscr{L} := \left\{ \frac{(\mu/n)^{\alpha-\delta} - 1}{\alpha-\delta} - \frac{(1-\mu/n)^{\alpha+\delta} - 1}{\alpha+\delta} : \alpha, \delta \in R \right\},$$

where logit link is given by $g_0(\mu) = \lim_{\alpha,\delta\to 0} g(\mu;\alpha,\delta)$, which can be seen by applying L'Hôspital's rule twice. In Preigbon family, α corresponds to the symmetry of the distribution of latent variable associated with a given model from $\mathscr{L}$ and the heaviness of tails of this distributions is parametrized by δ. In this case the carriers are $w_1 = 0.5(\log^2(\hat{\mu}/n) - \log^2(1\hat{\mu}/n))$ and $w_2 = -0.5(\log^2(\hat{\mu}/n) + \log^2(1\hat{\mu}/n))$. The comprehensive list of other families can be found in Stukel (1988). The carriers used in our tests are summarized in Table 1.

4.2 Procedure for finding the order in classifier chains

The methodology described in the previous section can be used to determine the optimal ordering in LCC. We propose to use a forward procedure such that at each stage we choose the label corresponding to best specified logistic model. The single step can be characterized

in the following way. We fit logistic models in which features $\mathbf{x}$ and labels found in the previous steps are treated as input variables whereas the candidate labels are treated as response variables. We select the label corresponding to the best specified model. The overall scheme is shown in Algorithm 1.

Algorithm 1: Pseudo-code of the procedure for finding the order in classifier chains.

Data: training set $\mathscr{D} = (\mathbf{x}^{(i)}, \mathbf{y}^{(i)})$, $i = 1, \ldots, n$
Initialize:
$\pi := \emptyset$
$I := (1, \ldots, K_n)$
$\mathbf{z}_{\text{act}} := \mathbf{x}$
for $k \leftarrow 1$ **to** K_n **do**
  $\hat{k} = \arg\min_{k \in I} D_{\hat{w}}(y_k, \mathbf{z}_{\text{act}})$, where D is defined in (12)
  $\mathbf{z}_{\text{act}} \leftarrow \mathbf{z}_{\text{act}} \cup \{y_{\hat{k}}\}$
  $I \leftarrow I \setminus \{\hat{k}\}$
  $\pi \leftarrow \pi \cup \{\hat{k}\}$
Output: ordering of labels π

In the following, we will give the heuristic justification of the above procedure. Consider the general situation in which we would like to test whether the given variable w is significant in the fitted logistic model. Assume that $P(y = 1|\mathbf{z}, w) = \sigma(\mathbf{z}'\theta + w\gamma)$, where $(\mathbf{z}', w)'$ is a vector of explanatory variables and θ, γ are unknown parameters. Let

$$D_w(y, \mathbf{z}) := 2\{l_{y,(\mathbf{z},w)}(\hat{\theta}_e, \hat{\gamma}_e) - l_{y,\mathbf{z}}(\hat{\theta})\} \tag{13}$$

be a deviance statistic, where $\hat{\theta}$ is an estimator based on smaller model (with variable w omitted), whereas $\hat{\theta}_e, \hat{\gamma}_e$ denote estimators based on extended model with an additional variable w. Let $D_{\gamma=0}$ denotes deviance (13) corresponding to the case $\gamma = 0$ and $D_{\gamma\neq 0}$ be the analogous quantity, corresponding to the case $\gamma \neq 0$.

Theorem 2 *Assume that:*

1. $E||\mathbf{z}||^2 < \infty$,
2. *there exists constant* c_1 *such that*

$$P\left[\lambda_{\min}[H_{y,(\mathbf{z}',w)'}(\theta, \gamma)/n] \geq c_1\right] \to 1$$

($\lambda_{\min}$ *is a minimal eigenvalue).*

Then, $P[D_{\gamma\neq 0} > D_{\gamma=0}] \to 1$*, as* $n \to \infty$.

For the proof, see in Section 8.3. Part of the proof concerning the asymptotic performance of $D_{\gamma=0}$ follows from Fahrmeir (1987), whereas the asymptotic convergence of $D_{\gamma\neq 0}$, under the above conditions, was proved in Teisseyre (2013). Observe that assumptions 1, 2 in above Theorem are analogous to assumptions 2, 4 in Theorem 1. The above Theorem is formulated for the case when there is only one additional variable w, but it can be obviously expanded to the case of more additional variables.

The above Theorem aims to justify a procedure described in Algorithm 1. Namely, deviance $D_{\gamma=0}$ corresponds to the correct specification, whereas $D_{\gamma\neq 0}$ corresponds to the incorrect specification. Note however, that there are small differences between situation described

in Theorem 2 and setting described in Section 4.1 (and applied in Algorithm 1). First, observe that the extended model in (12) is not exactly a true one (as in (13)), because it is based on approximation given in (11). Secondly, observe that an additional variable $\hat{w}$ in (12) depends on the estimator $\hat{\theta}$ calculated for the smaller model, which is not the case in (13).

5 Inference in classifier chain model

The classifier chain model allows to estimate the joint probability distribution. Having estimated the joint distributions, one would usually like to infer from this distribution, i.e. to make a prediction for some new instance $\mathbf{x}$. The Bayes optimal rule is the mode of the joint distribution $\mathbf{y}^*(\mathbf{x}) = \arg\max_{y\in\{0,1\}^{K_n}} P(\mathbf{y}|\mathbf{x})$, which is estimated by $\hat{\mathbf{y}}(\mathbf{x}) = \arg\max_{\mathbf{y}\in\{0,1\}^{K_n}} \hat{P}(\mathbf{y}|\mathbf{x})$, based on training set. In Theorem 1 we have shown that for classifier chain model, $\hat{\mathbf{y}}(\mathbf{x})$ converges to $\mathbf{y}^*(\mathbf{x})$ in probability, even in the case of large number of labels, provided that the "true" ordering of labels is known and certain assumptions on the design matrix and number of labels are satisfied. However, from the practical point of view, the problematic is computation of $\hat{\mathbf{y}}(\mathbf{x})$ as it requires evaluation of 2^{K_n} possible labellings for one new instance $\mathbf{x}$. This approach is often referred to as "exhaustive inference". The "exhaustive inference" is limited to data sets with a small to moderate number of labels, say, not more than about 15. In "multi-label community", the method which combines the classifier chains with "exhaustive inference" is called Probabilistic classifier chains (PCC), see in Dembczyński et al. (2010) and Read et al. (2011). Chain structure of the model, suggests to use "greedy inference" by successively choosing the most probable label according to each of the classifiers predictions, i.e. when the assumed order of labels is π^0, then we apply the following procedure

- find $\hat{y}_{\pi^0(1)} = \arg\max_{y\in\{0,1\}} \hat{P}(y_{\pi^0(1)} = y|\mathbf{x})$,
- find $\hat{y}_{\pi^0(2)} = \arg\max_{y\in\{0,1\}} \hat{P}(y_{\pi^0(2)} = y|\hat{y}_{\pi^0(1)}, \mathbf{x})$,
- ...
- find $\hat{y}_{\pi^0(K_n)} = \arg\max_{y\in\{0,1\}} \hat{P}(y_{\pi^0(K_n)} = y|\hat{y}_{\pi^0(1)}, \ldots, \hat{y}_{\pi^0(K_n-1)}, \mathbf{x})$,

which requires K_n operations instead of 2^{K_n}. This approach, introduced originally in Read et al. (2009), is usually simply referred to as classifier chains (CC). Since the naming can be a little confusing, we stress that the learning stage is the same in both PCC and CC, whereas the difference is in the inference stage. Note that the "greedy inference" does not guarantee that the mode will be identified correctly, even when the true joint distribution is known. Consider an example with two binary labels y_1 and y_2 and the ordering $\pi^0 = (1,2)$. Let $P(y_1 = 1|\mathbf{x}) = 0.6$, $P(y_1 = 0|\mathbf{x}) = 0.4$, $P(y_2 = 1|\mathbf{x}, y_1 = 1) = P(y_2 = 0|\mathbf{x}, y_1 = 1) = 0.5$ and $P(y_2 = 1|\mathbf{x}, y_1 = 0) = 0.9$. Then the "greedy inference" yields labelling $\mathbf{y} = (1,1)'$ (or $\mathbf{y} = (1,0)'$), whereas the true mode is $\mathbf{y}^* = (0,1)'$. The other problem with "greedy inference", addressed in Senge et al. (2012), is an error propagation along the chain. Since each classifier relies on its predecessors, a false prediction might be propagated. In a consequence, the labels which are placed at the end of the chain, carry the highest errors.

The problem of inference from the joint distribution can be visualise using the binary tree, in which nodes represent labels, edges represent conditional probabilities and leaves represent particular labellings. Figure 3 shows an example tree for two labels. The inference task is equivalent to finding the optimal path in a rooted, complete binary tree of height K_n. The "exhaustive inference" requires evaluation of all 2^{K_n} paths in the tree, on the other hand, the "greedy inference" corresponds to one chosen path.

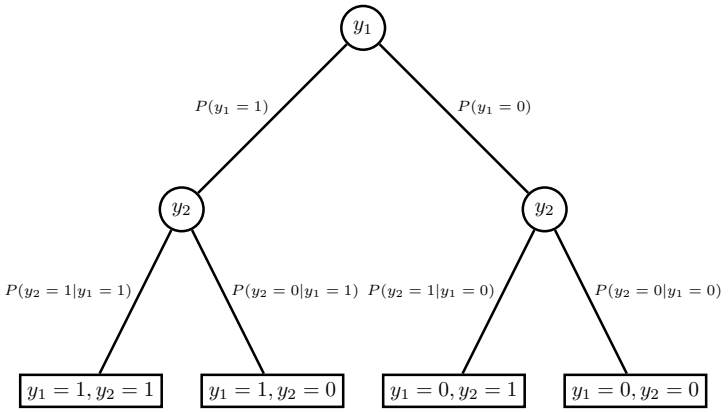

Fig. 3 Binary tree used for inference with two labels.

To overcome the problem of time complexity of "exhaustive inference", from one side and the possible lower accuracy of "greedy inference", from the other side, some authors develop approximate inference schemes that trade off accuracy against efficiency in a reasonable way. For example, Kumar et al. (2013) proposed using beam search algorithm for inference to explore the binary tree. The idea is to keep b candidate solutions at each level of the tree, where b is a user-defined parameter known as the beam width, which represent the best partial solutions seen thus far. Then the tree is explored in a breadth-first fashion using these solutions. "Beam search inference" requires bK operations and allows to explore more paths in the binary tree, than for "greed inference". Of course, as in the case of "greedy inference", this approach may also not result in finding the optimal label vector. Observe that "Beam search inference" with $b = 1$ is equivalent to "greedy inference", whereas $b = \infty$ is equivalent to "exhaustive inference". The other simple solution is to limit the inference only to the label combinations that appear in the training set. Since the inference problem is not a main issue of this paper, we limit our empirical experiments to two basic methods: "exhaustive inference" and "greedy inference".

6 Empirical evaluation

6.1 Performance of goodness of specification measures

The aim of the first experiment is to compare the goodness of specification tests described in Section 4.1. We use 10 models generated in the following way. We first generate feature vector $\mathbf{x}$ from uniform distribution on $[-4,4]$ and then the consecutive labels $y_1,\ldots,y_{K_n}$ are generated using classifier chain model (3) according to the order $\pi^*(k) = k$. The parameters corresponding to the considered models are summarized in Table 2. The first two simple models match the one discussed at the beginning of Section 4, with $a = 3$ and $a = 5$, respectively. The structure of the models was chosen so as to investigate what is the influence of the: number of labels, number of features, values of parameters corresponding to features and values of parameters corresponding to labels. We use the procedure described in 4.2 in combination with methods summarized in Table 1. For each $n = 50, 100, \ldots$, we repeat 200 times: generate data with n instances and apply the procedure for finding the optimal

Table 2 Characteristics of the considered models.

Model	Parameters	Model	Parameters
M1	$\theta_1 = (0,1)'$ $\theta_2 = (0,1,3)'$	M7	$\theta_1 = (2,-2,1)'$ $\theta_2 = (2,-2,1,5)'$ $\theta_3 = (2,-2,1,5,-5)'$ $\theta_4 = (2,-2,1,-5,5,-5)'$
M2	$\theta_1 = (0,1)'$ $\theta_2 = (0,1,5)'$	M8	$\theta_1 = (2,-2,1)'$ $\theta_2 = (2,-2,1,2)'$ $\theta_3 = (2,-2,1,2,-2)'$ $\theta_4 = (2,-2,1,-2,2,-2)'$
M3	$\theta_1 = (2,-2,1)'$ $\theta_2 = (2,-2,1,5)'$ $\theta_3 = (2,-2,1,5,-5)'$ $\theta_4 = (2,-2,1,-5,5,-5)'$ $\theta_5 = (2,-2,1,5,-5,5,-5)'$ $\theta_6 = (2,-2,1,5,-5,5,-5,5)'$	M9	$\theta_1 = (2,-2,1)'$ $\theta_2 = (2,-2,1,10)'$ $\theta_3 = (2,-2,1,10,-10)'$ $\theta_4 = (2,-2,1,-10,10,-10)'$
M4	$\theta_1 = (2,-2,1)'$ $\theta_2 = (2,-2,1,5)'$ $\theta_3 = (2,-2,1,5,-5)'$ $\theta_4 = (2,-2,1,-5,5,-5)'$ $\theta_5 = (2,-2,1,5,-5,5,-5)'$	M10	$\theta_1 = (5,-5,2)'$ $\theta_2 = (5,-5,2,5)'$ $\theta_3 = (5,-5,2,5,-5)'$ $\theta_4 = (5,-5,2,-5,5,-5)'$
M5	$\theta_1 = \mathbf{a}$ $\theta_2 = (\mathbf{a}',5)'$ $\theta_3 = (\mathbf{a}',5,-5)'$ $\theta_4 = (\mathbf{a}',-5,5,-5)'$ $\mathbf{a} = (1,-1,1,-1,1,-1,1,-1,1,-1)'$	M11	$\theta_1 = \mathbf{a}$ $\theta_2 = (\mathbf{a}',-8)'$ $\theta_3 = (\mathbf{a}',1,3)'$ $\theta_4 = (\mathbf{a}',0.5,5,10)'$ $\mathbf{a} = (1,-1,1,-1,1,-1,1,-1,1,-1)'$
M6	$\theta_1 = (1,-3,0.5)'$ $\theta_2 = (1.5,-2.5,1,5)'$ $\theta_3 = (2,-2,1.5,5,-5)'$ $\theta_4 = (2.5,-1.5,2,-5,,5-5)'$	M12	$\theta_1 = \mathbf{a}$ $\theta_2 = (\mathbf{a}',5)'$ $\theta_3 = (\mathbf{a}',5,-5,)'$ $\theta_4 = (\mathbf{a}',-5,5,-5)'$ $\theta_5 = (\mathbf{a}',5,-5,5,-5)'$ $\theta_6 = (\mathbf{a}',5,-5,5,-5,5)'$ $\theta_7 = (\mathbf{a}',5,-5,5,-5,5,-5)'$ $\theta_8 = (\mathbf{a}',5,-5,5,-5,5,-5,5)'$ $\theta_9 = (\mathbf{a}',5,-5,5,-5,5,-5,5,-5)'$ $\theta_{10} = (\mathbf{a}',5,-5,5,-5,5,-5,5,-5,5)'$ $\mathbf{a} = (1,-1,1,-1,1,-1,1,-1,1,-1)'$

ordering. This enables to estimate the probability of choosing the correct ordering. Figure 4 shows the results. The dashed line corresponds to the random choice of the optimal ordering. We also tested using minus log-likelihood function (this measure was used in Kumar et al. (2013)) instead of $D(y_k, \mathbf{z}_{\text{act}})$ in Algorithm 1. This gives very poor results (comparable with the random choice), so we do not present the results in the figures. The structure of the model affects the probability of choosing the correct ordering. Observe that, in all cases, the results are significantly better than the baseline (random choice). The method proposed by Morgan performs poorly in the majority of cases. The methods, which use two carriers (Preigbon, Stukel, Prentice), usually outperform those which use only one carrier. This is clearly seen for models (M4) and (M6). This phenomenon can be explained by the fact that the methods from the first group are able to represent wider family of possibly misspecified models. In the case of model (M12), the probabilities are relatively low (about 0.01 for $n = 4000$ and Preigbon method), but they are still much larger than the random choice, which is $1/K_n! \approx 10^{-7}$. On the other hand, in this case, the use of Preigbon methods improves the probability of correct mode selection significantly (see in the next section).

Note also that the probability of choosing the correct ordering decreases when the number of labels increases (compare models (M3) and (M4) or (M5) and (M10)) and the probability of choosing the correct ordering decreases when the absolute values of parameters corresponding to labels decrease (compare models (M8) and (M9)).

Table 3 Basic statistics for the benchmark datasets.

Dataset	Domain	#observations	#features	#labels
cal500	music	502	68	174
emotions	music	593	72	6
flags	images (toy)	194	19	7
mediamill	video	10000	120	101
scene	images	2407	294	6
yeast	biology	2417	103	14

6.2 Consistency of the joint mode selection

In the second experiment, we illustrate the theoretical result from Theorem 1 concerning the consistency of the joint mode selection. In this experiment we use the same models (M1)-(M12) as in the previous section. The models are generated as was described previously and additionally we generate test set containing 200 observations. For each observation in the test set we check whether the correct mode was selected. The whole procedure is repeated 200 times, which yields the estimates of the probability of correct mode selection. As the correct ordering is unknown in practical applications, we use the procedure described in Algorithm 1, combined with Preigbon method. Figure 5 shows probabilities of correct mode selection with respect to the number of observations in the training set for the correct ordering, selected ordering and wrong ordering (reversed correct order). To assess the direct effect of ordering on the joint mode estimation, we use "exhaustive inference" in the case of all models, except model (M12), for which "greedy inference" was used, due to computational costs. First of all, it is seen that the ordering of labels may affect the probability of correct mode selection, although for some models (e.g. (M3) and (M4)) the differences are not significant. Secondly, it is seen that Algorithm 1, combined with Preigbon method, performs well in practice- the results are very close to those for correct ordering (for models (M1), (M2), (M5) they practically coincide). In the case of model (M12), the differences between the selected ordering and the wrong one are significant, although it is difficult to determine the true ordering exactly (compare Figure 4, model (M12)).

6.3 Experiments on benchmark data sets

To investigate the performance of the proposed approach, we also carried out experiments on benchmark datasets. The datasets are publicly available at `http://mulan.sourceforge.net/datasets-mlc.html` (at this website one can also find the detailed description of each data set, as well as the references). The details of the data sets are summarized in Table 3. We compare the following methods.

- (BR) Binary relevance method, in which we build K_n separate classifiers for each label. Logistic model is used as a single classifier.
- (CC EX) Classifier chain with the logistic model as a single classifier and "exhaustive inference". The order of fitting models in the chain corresponds to the original ordering of labels in dataset.
- (CC PREIGBON EX) Classifier chain with the logistic model as a single classifier and "exhaustive inference". The order of fitting models in the chain is determined using Algorithm 1 combined with Preigbon method.

– (CC GR) Classifier chain with the logistic model as a single classifier and "greedy inference". The order of fitting models in the chain corresponds to the original ordering of labels in dataset.
– (CC PREIGBON GR) Classifier chain with the logistic model as a single classifier and "greedy inference". The order of fitting models in the chain is determined using Algorithm 1 combined with Preigbon method.

All above methods have been implemented in R system (R Core Team (2013)), for the purpose of this paper. The common part of all considered methods is a logistic model. To reduce the variance of the estimators, we used regularized versions of maximum likelihood estimators, with l_2 penalty and small value of penalty parameter $\lambda = 0.001$. Let $\mathbf{y}$ be the vector of true labels and $\hat{\mathbf{y}}$ be the predicted vector of labels for given test instance. We consider the following evaluation measures

$$\text{Hamming}(\mathbf{y},\hat{\mathbf{y}}) = \frac{1}{K_n}\sum_{k=1}^{K_n}\mathbf{I}[y_k = \hat{y}_k],$$

$$\text{Subset accuracy}(\mathbf{y},\hat{\mathbf{y}}) = \mathbf{I}[\mathbf{y} = \hat{\mathbf{y}}],$$

$$\text{Recall}(\mathbf{y},\hat{\mathbf{y}}) = \frac{|\{k : y_k = 1, \hat{y}_k = 1\}|}{|\{k : y_k = 1\}|},$$

$$\text{Precision}(\mathbf{y},\hat{\mathbf{y}}) = \frac{|\{k : y_k = 1, \hat{y}_k = 1\}|}{|\{k : \hat{y}_k = 1\}|},$$

$$\text{F measure}(\mathbf{y},\hat{\mathbf{y}}) = 2 \cdot \frac{\text{Precision}(\mathbf{y},\hat{\mathbf{y}}) \cdot \text{Recall}(\mathbf{y},\hat{\mathbf{y}})}{\text{Precision}(\mathbf{y},\hat{\mathbf{y}}) + \text{Precision}(\mathbf{y},\hat{\mathbf{y}})}.$$

The subset accuracy is maximized by the methods which allow to find the mode of the joint distribution (e.g. classifier chains). On the other hand, the Hamming measure is maximized by the methods which find the marginal modes (BR). The marginal and joint modes coincide when the labels are conditionally independent given feature vector $\mathbf{x}$, but this is usually not the case. For further discussion on different evaluation measures and loss functions see in Dembczyński et al. (2012). Since in many multi-label problems we are interested in predicting the presence of certain properties of the objects (the presence of k-th property is usually coded as $y_k = 1$), it is worthwhile to consider additional measures: Recall, Precision and F-measure. For example, in text categorization, $y_k = 1$ indicates that the given text has been assigned to k-th topic. Recall indicates how many labels are correctly predicted as 1 among those equal 1, precision shows how many labels are correctly predicted as 1 among those predicted as 1 and F-measure is a harmonic mean between recall and precision.

To assess the considered methods, we apply 5-fold cross-validation. In each cross-validation loop, the models are built on the four, train folds and the prediction is made on the remaining, test fold. The above evaluation measures are averaged over observations in the test fold. Finally, the reported results are averaged over all cross-validation folds.

We performed two experiments. In the first one, we consider datasets with small or moderate number of labels: emotions, flags and scene. For the remaining datasets (yeast, mediamill, cal500) we have limited the number of labels to 10 in such a way that we keep the most frequent labels (and remove all instances having only relevant or only irrelevant labels). This approach was also used in Dembczyński et al. (2010) to test PCC method. The limited number of labels allows to use "exhaustive inference". Thus we eliminate the effect associated with the type of prediction method and we can directly assess how the ordering of labels affects the results. The results for the first experiment are shown in Tables 4-9. The

bold font corresponds to the maximal value in the given column. It is notable that, application of Algorithm 1, combined with Preigbon method, improves the subset accuracy, relative to original ordering of labels for 5 out of 6 datasets. This occurs in both cases: for exhaustive and greedy inferences. Moreover, in the majority of cases, CC PREIGBON EX outperforms other methods with respect to all measures except Hamming measure. The "exhaustive inference" outperforms "greedy inference" which is obvious as in the latter case we explore the limited number of possible labellings. We have the highest values of Hamming measure for BR method, which is consistent with observations of other authors and also confirms the theoretical results concerning Hamming measure, given in Dembczyński et al. (2012).

The second experiment was performed for datasets: yeast, mediamill, cal500, with the original number of labels (which varies from 14 to 174). For this setting, we only used 3 methods: BR, CC GR and CC GR PREIGBON, as other methods are infeasible for such a large number of labels. The results for the second experiment are shown in Tables 10-12. It is seen that also in this case the application of Algorithm 1, combined with Preigbon method, improves the subset accuracy and recall for all three datasets and improves the F-measure for two datasets.

Table 4 Results for emotions data set.

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	**0.7917** (0.0062)	0.2478 (0.0092)	0.6276 (0.015)	0.6456 (0.0174)	0.6027 (0.0125)
CC EX	0.7791 (0.0031)	0.2917 (0.011)	0.6619 (0.0124)	0.6416 (0.0106)	0.6247 (0.0084)
CC PREIGBON EX	0.7889 (0.0036)	**0.3052** (0.0083)	**0.6844** (0.0128)	**0.6627** (0.0103)	**0.6448** (0.0072)
CC GR	0.776 (0.0051)	0.2731 (0.0094)	0.6431 (0.0164)	0.6405 (0.015)	0.6126 (0.0139)
CC PREIGBON GR	0.7839 (0.0051)	0.29 (0.0099)	0.6661 (0.0148)	0.6531 (0.0104)	0.6303 (0.0108)

Table 5 Results for scene data set.

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	**0.8948** (9e-04)	0.5272 (0.0055)	0.6597 (0.0032)	0.6217 (0.0043)	0.6279 (0.0033)
CC EX	0.8912 (0.0027)	0.6315 (0.0056)	0.7035 (0.008)	0.706 (0.0085)	0.6967 (0.0078)
CC PREIGBON EX	0.8945 (0.0023)	**0.6386** (0.0069)	**0.7145** (0.006)	**0.7121** (0.0069)	**0.7051** (0.0063)
CC GR	0.8905 (0.0032)	0.6278 (0.0079)	0.701 (0.0104)	0.7021 (0.0102)	0.6934 (0.0098)
CC PREIGBON GR	0.8922 (0.0028)	0.6298 (0.0085)	0.7062 (0.0078)	0.7034 (0.0089)	0.6966 (0.0082)

Table 6 Results for flags data set.

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	**0.7297** (0.0124)	0.139 (0.0098)	0.6757 (0.0191)	**0.695** (0.0124)	0.6723 (0.0117)
CC EX	0.7173 (0.0126)	0.237 (0.0316)	**0.6785** (0.0154)	0.6733 (0.0124)	0.6736 (0.0134)
CC PREIGBON EX	0.7179 (0.0056)	**0.2625** (0.029)	0.6749 (0.0138)	0.6774 (0.0096)	**0.6741** (0.0111)
CC GR	0.6982 (0.0118)	0.2012 (0.0156)	0.6559 (0.0199)	0.6535 (0.0141)	0.652 (0.0168)
CC PREIGBON GR	0.7144 (0.0145)	0.2421 (0.0408)	0.6733 (0.0166)	0.6771 (0.0088)	0.673 (0.0131)

Table 7 Results for yeast data set (top 10 labels).

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	**0.7461** (0.004)	0.1589 (0.0077)	0.6175 (0.0042)	**0.6972** (0.0076)	0.626 (0.0052)
CC EX	0.7352 (0.0044)	**0.2453** (0.008)	**0.6572** (0.0074)	0.6578 (0.0066)	**0.6345** (0.0069)
CC PREIGBON EX	0.7328 (0.0061)	0.2379 (0.0098)	0.6499 (0.0095)	0.661 (0.0077)	0.631 (0.0076)
CC GR	0.7255 (0.0047)	0.2214 (0.0057)	0.6208 (0.0063)	0.6428 (0.0083)	0.6056 (0.007)
CC PREIGBON GR	0.7313 (0.006)	0.2143 (0.0099)	0.6343 (0.0094)	0.6601 (0.0086)	0.6202 (0.0072)

Table 8 Results for cal500 data set (top 10 labels).

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	0.6052 (0.0075)	0.004 (0.0024)	0.7193 (0.0075)	0.6626 (0.0081)	0.6713 (0.007)
CC EX	0.6016 (0.0087)	0.0159 (0.0051)	0.7065 (0.0115)	0.661 (0.0087)	0.6649 (0.0091)
CC PREIGBON EX	0.6044 (0.0085)	**0.016** (0.0051)	0.7121 (0.0106)	0.6631 (0.0085)	0.6679 (0.0089)
CC GR	0.5954 (0.0091)	0.0139 (0.0074)	0.6998 (0.009)	0.6552 (0.0098)	0.6579 (0.0076)
CC PREIGBON GR	0.608 (0.0083)	0.01 (0.0055)	**0.7141** (0.0082)	**0.6679** (0.0096)	0.6722 (0.0075)

Table 9 Results for mediamill data set (top 10 labels).

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	**0.8277** (0.0013)	0.1642 (0.0052)	0.5844 (0.0035)	**0.7456** (0.0012)	**0.6246** (0.0026)
CC EX	0.8165 (0.0023)	0.2152 (0.0058)	0.6049 (0.0066)	0.6836 (0.0049)	0.6094 (0.0054)
CC PREIGBON EX	0.8186 (0.0021)	**0.2168** (0.007)	0.6033 (0.0075)	0.6873 (0.0035)	0.6103 (0.0057)
CC GR	0.8199 (9e-04)	0.2046 (0.004)	0.6007 (0.005)	0.703 (0.0017)	0.6143 (0.0034)
CC PREIGBON GR	0.8211 (0.0025)	0.2098 (0.0048)	0.5772 (0.0083)	0.7087 (0.0052)	0.6032 (0.0055)

Table 10 Results for yeast data set (all labels).

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	**0.7959** (0.0032)	0.1411 (0.0061)	0.5871 (0.0052)	**0.6963** (0.0075)	**0.6083** (0.0051)
CC GR	0.7708 (0.0102)	0.175 (0.0265)	0.6033 (0.0091)	0.6221 (0.0156)	0.5867 (0.0059)
CC PREIGBON GR	0.7808 (0.0044)	**0.1812** (0.0052)	**0.6059** (0.0117)	0.6526 (0.007)	0.6017 (0.0077)

Table 11 Results for cal500 data set (all labels).

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	0.7676 (0.0037)	0 (0)	0.3339 (0.008)	0.2719 (0.0045)	0.2937 (0.0033)
CC GR	**0.7825** (0.0042)	0 (0)	0.3431 (0.0109)	**0.3017** (0.0087)	**0.3146** (0.0082)
CC PREIGBON GR	0.7773 (0.0074)	0 (0)	**0.3454** (0.0131)	0.295 (0.0088)	0.3109 (0.0048)

Table 12 Results for mediamill data set (all labels).

method	Hamming	Subset accuracy	Recall	Precision	F-measure
BR	**0.9441** (0.0046)	0 (0)	0.4538 (0.0045)	**0.3749** (0.0404)	**0.3837** (0.0205)
CC GR	0.8989 (0.0026)	0 (0)	0.4632 (0.0097)	0.2175 (0.0063)	0.2769 (0.0043)
CC PREIGBON GR	0.9118 (0.0068)	**0.0028** (0.0016)	**0.4887** (0.0128)	0.2744 (0.0276)	0.3191 (0.0182)

7 Conclusions

In this paper we studied the large sample properties of the logistic classifier chain (LCC) model. We found the conditions under which the estimated joint mode is consistent. In particular it follows from Theorem 1 that the number of labels should be relatively small compared with the number of training examples. Secondly, we have shown that when the ordering of building models is incorrect, the consecutive probabilities in the chain are approx-

imated by logistic models as closely as possible with respect to Kullback-Leibler distance. We have shown that the ordering of labels may affect probability of joint mode selection significantly. The proposed procedure of ordering the labels, based on measures of correct specification, performs well in practice, which is confirmed by experiments on artificial and real datasets.

Acknowledgements

Research of Paweł Teisseyre was supported by the European Union from resources of the European Social Fund within project 'Information technologies: research and their interdisciplinary applications' POKL.04.01.01-00-051/10-00.

8 Proofs

8.1 Proof of Theorem 1

We present the proof for $K_n \to \infty$; the proof for the constant K_n goes along the same lines, but requires some simple modifications. For simplicity, we will write l_k, s_k and H_k instead of $l_{y_k,\mathbf{z}_k}$, $s_{y_k,\mathbf{z}_k}$, $H_{y_k,\mathbf{z}_k}$, respectively. We will show that $P[\hat{\mathbf{y}}(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})] \to 0$, as $n \to \infty$. Recall that $\mathbf{z}_k = (\mathbf{x}', \mathbf{y}'_{1:(k-1)})'$. Using Lemmas 2 and 3 we have

$$
\begin{aligned}
&P[\hat{\mathbf{y}}(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})] \leq P\left[\max_{\mathbf{y}\in\{0,1\}^{K_n}} |\hat{P}(\mathbf{y}|\mathbf{x}) - P(\mathbf{y}|\mathbf{x})| > \frac{\varepsilon}{2}\right] \leq \\
&P\left[\max_{\mathbf{y}\in\{0,1\}^{K_n}} \sum_{k=1}^{K_n} \left|\sigma(\mathbf{z}_k'\hat{\theta}_k)^{y_k}[1-\sigma(\mathbf{z}_k'\hat{\theta}_k)]^{1-y_k} - \sigma(\mathbf{z}_k'\theta_k)^{y_k}[1-\sigma(\mathbf{z}_k'\theta_k)]^{1-y_k}\right| > \frac{\varepsilon}{2}\right] \leq \\
&P\left[\max_{\mathbf{y}\in\{0,1\}^{K_n}} \max_{1\leq k\leq K_n} \left|\sigma(\mathbf{z}_k'\hat{\theta}_k)^{y_k}[1-\sigma(\mathbf{z}_k'\hat{\theta}_k)]^{1-y_k} - \sigma(\mathbf{z}_k'\theta_k)^{y_k}[1-\sigma(\mathbf{z}_k'\theta_k)]^{1-y_k}\right| > \frac{\varepsilon}{2K_n}\right] \leq \\
&\sum_{k=1}^{K_n} P\left[\max_{\mathbf{y}\in\{0,1\}^{K_n}} \left|\sigma(\mathbf{z}_k'\hat{\theta}_k) - \sigma(\mathbf{z}_k'\theta_k)\right| > \frac{\varepsilon}{2K_n}\right]. \qquad (14)
\end{aligned}
$$

Using mean value theorem and Schwarz inequality the probability under the above sum can be bounded by

$$
P\left[\max_{\mathbf{y}\in\{0,1\}^K} |\mathbf{z}_k'\hat{\theta}_k - \mathbf{z}_k'\theta_k| > \frac{\varepsilon}{2K_n}\right] \leq P\left[||\hat{\theta}_k - \theta_k|| > \frac{\varepsilon}{2K_n^{3/2}}\right],
$$

where the last inequality follows from the fact that $||\mathbf{z}_k||^2 \leq 4K_n$, for sufficiently large n (as first p coordinates of $\mathbf{z}_k$ corresponds to a fixed point $\mathbf{x}$). Denote by $\varepsilon_n := \frac{\varepsilon}{2K_n^{3/2}}$. Define $\theta_{k,u} := \theta_k + u\varepsilon_n/2$, for $u \in R^p$ such that $||u|| = 1$. Using Taylor expansion and concavity of l_k, we have for some point θ_k^* between $\theta_{k,u}$ and θ_k

$$
\begin{aligned}
&P\left[||\hat{\theta}_k - \theta_k|| > \varepsilon_n\right] \leq P\left[\exists u : ||u|| = 1, l_k(\theta_{k,u}) - l_k(\theta_k) > 0\right] = \\
&P\left[\exists u : ||u|| = 1, u's_k(\theta_k) > \varepsilon_n u'H_k(\theta_k^*)u/4\right]. \qquad (15)
\end{aligned}
$$

Using Schwarz inequality we have

$$
\max_{1\leq k\leq K_n} \max_{1\leq i\leq n} |\mathbf{z}^{(i)'}(\theta_k - \theta_k^*)|^2 \leq \max_{1\leq i\leq n} ||\mathbf{z}_{K_n}^{(i)}||^2\varepsilon_n^2 \leq \max_{1\leq i\leq n} ||\mathbf{x}_{K_n}^{(i)}||^2\varepsilon_n^2 + K_n\varepsilon_n^2 \xrightarrow{P} 0, \quad (16)
$$

where the last convergence follows from Lemma 5, Markov inequality and assumption 2. In view of (16), the assertion of Lemma 4 is satisfied for all $1 \leq k \leq K_n$ and thus the probability in (15) can be bounded from above by

$$
\begin{aligned}
&P\left[\exists u : ||u|| = 1, u's_k(\theta_k) > \varepsilon_n cu'H_k(\theta_k)u/4\right] \leq \\
&P\left[||s_k(\theta_k)|| > \varepsilon_n c\lambda_{\min}[H_k(\theta_k)]/4\right] \leq P\left[||s_k(\theta_k)|| > \varepsilon_n cc_1 n/4\right], \qquad (17)
\end{aligned}
$$

where the second inequality follows from assumption 4. Let $r_n := p + K_n - 1$. Using Lemma 1 and the fact that we have (for some constant c_2, with probability tending to 1)

$$
tr(Z_k'Z_k) \leq r_n c_2 n,
$$

the probability in (17) can be further bounded by

$$P\left[||s_k(\theta_k)||^2 > \frac{c^2c_1^2\varepsilon_n^2n^2r_nc_2}{16r_nc_2}\right] \leq P\left[||s_k(\theta_k)||^2 > \frac{c^2c_1^2}{16c_2}\frac{\varepsilon_n^2n}{r_n}tr(Z_k'Z'k)\right] \leq$$
$$\exp\left[-\frac{1}{20}\frac{c^2c_1^2}{16c_2}\frac{\varepsilon_n^2n}{r_n}\right]. \tag{18}$$

It follows from (18) and assumption 3 that probability in (14) converges to zero, which ends the proof. □

8.2 Proof of Corollary 1

The probability in Corollary 1 can be bounded from above as

$$P(||\hat{\theta}-\theta||^2 > \varepsilon K_n^{-2}) = P\left[\sum_{k=1}^{K_n}||\hat{\theta}_k-\theta_k||^2 > \varepsilon K_n^{-2}\right] \leq P(\max_{1\leq k\leq K_n}||\hat{\theta}_k-\theta_k||^2 > \varepsilon K_n^{-3}) \leq$$
$$\sum_{k=1}^{K_n}P(||\hat{\theta}_k-\theta_k|| > \varepsilon K_n^{-3/2}) \to 0,$$

from the proof of Theorem 1 (see (15)-(18)). □

8.3 Proof of Theorem 2

Consider first term $D_{\gamma=0}$. It follows from Lemma 4.2.7 and Remark 4.2.3 in Teisseyre (2013), that assumptions of Theorem 2 imply assumptions of Theorem 1 in Fahrmeir (1987), which states that $D_{\gamma=0} \xrightarrow{d} \chi_1^2$, where $\xrightarrow{d}$ denotes convergence in distribution and χ_1^2 is a random variable distributed according to chi-squared distribution with 1 degree of freedom. Consider term $D_{\gamma\neq 0}$. It follows from Proposition 4.2.10 in Teisseyre (2013), that under assumptions of Theorem 2, $P[D_{\gamma\neq 0} > w_n] \to 1$, as $n \to \infty$, for some sequence $w_n \to \infty$. This implies the assertion. □

8.4 Auxiliary facts

Lemma 2 *Assume that* $P(\mathbf{y}^*(\mathbf{x})|\mathbf{x}) > P(\mathbf{y}|\mathbf{x}) + \varepsilon$, *for all* $\mathbf{y} \neq \mathbf{y}^*(\mathbf{x})$ *and some* $\varepsilon > 0$. *Then* $\hat{\mathbf{y}}(\mathbf{x}) \neq \mathbf{y}^*(\mathbf{x})$ *implies that*

$$\max_{\mathbf{y}\in\{0,1\}^K}|\hat{P}(\mathbf{y}|\mathbf{x}) - P(\mathbf{y}|\mathbf{x})| > \frac{\varepsilon}{2}.$$

Proof When $|P(\mathbf{y}^*(\mathbf{x})|\mathbf{x}) - \hat{P}(\mathbf{y}^*(\mathbf{x})|\mathbf{x})| > \frac{\varepsilon}{2}$ then the assertion holds. Consider the case $|P(\mathbf{y}^*(\mathbf{x})|\mathbf{x}) - \hat{P}(\mathbf{y}^*(\mathbf{x})|\mathbf{x})| \leq \frac{\varepsilon}{2}$. Then we have $\hat{P}(\hat{\mathbf{y}}(\mathbf{x})|\mathbf{x}) \geq \hat{P}(\mathbf{y}^*(\mathbf{x})|\mathbf{x})$ which implies $|P(\hat{\mathbf{y}}(\mathbf{x})|\mathbf{x}) - \hat{P}(\hat{\mathbf{y}}(\mathbf{x})|\mathbf{x})| > \frac{\varepsilon}{2}$ in view of the assumption. This ends the proof.

Lemma 3 *Let* $|a_k| \leq 1$, $|b_k| \leq 1$, *for* $k = 1, \ldots, K_n$. *Then*

$$\left|\prod_{k=1}^{K_n}a_k - \prod_{k=1}^{K_n}b_k\right| \leq \sum_{k=1}^{K_n}|a_k - b_k|.$$

Proof Using the following inequality

$$\left|\prod_{k=1}^{K_n} a_k - \prod_{k=1}^{K_n} b_k\right| = \left|\prod_{k=1}^{K_n} a_k - b_{K_n}\prod_{k=1}^{K_n-1} a_k + b_{K_n}\prod_{k=1}^{K_n-1} a_k - \prod_{k=1}^{K_n} b_k\right| =$$
$$\left|(a_{K_n} - b_{K_n})\prod_{k=1}^{K_n-1} a_k + b_{K_n}\left(\prod_{k=1}^{K_n-1} a_k - \prod_{k=1}^{K_n-1} b_k\right)\right| \leq |a_{K_n} - b_{K_n}| + \left|\prod_{k=1}^{K_n-1} a_k - \prod_{k=1}^{K_n-1} b_k\right|.$$

and induction, the assertion follows. □

Lemmas 4 and 5 below are proved in Mielniczuk and Teisseyre (2015).

Lemma 4 *If* $\max_{1\leq i\leq n} |\mathbf{z}^{(i)'}(\gamma-\theta)| \leq c$ *then for any vector* $u \in R^p$ *we have*

$$\exp(-3c)u'H_k(\theta)u \leq u'H_k(\gamma)u \leq \exp(3c)u'H_k(\theta)u.$$

Lemma 5 *The convergence* $\max_{1\leq i\leq n} ||\mathbf{x}^{(i)}||\varepsilon_n \xrightarrow{P} 0$ *is equivalent to* $||\mathbf{x}^{(n)}||\varepsilon_n \xrightarrow{P} 0$ *for non-decreasing sequence* $\varepsilon_n \to 0$.

References

J. Chen and Z. Chen. Extended BIC for small-n-large-P sparse glm. *Statistica Sinica*, 22: 555–574, 2012.

C. Czado and T. Santner. The effect of link misspecification on binary regression inference. *Journal of Statistical Planning and Inference*, 33:213–231, 1992.

K. Dembczyński, W. Cheng, and E. Hüllermeier. Bayes optimal multilabel classification via probabilistic classifier chains. In *Proceedings of the twenty-seventh international conference on machine learning*, volume 22, pages 109–117, 2010.

K. Dembczyński, W. Waegeman, W. Cheng, and E. Hüllermeier. *Machine Learning*, 88: 5–45, 2012.

L. Fahrmeir. Asymptotic testing theory for generalized linear models. *Statistics*, 1:65–76, 1987.

L. Fahrmeir and H. Kaufmann. Consistency and asymptotic normality of the maximum likelihood estimator in generalized linear models. *Annals of Statistics*, 1(13):342–368, 1985.

N. Hjort. Logistic regression when the model is wrong. Appendix to Statistical models for the probability of finding oil or gas. Norwegian Computing Centre Report., 1988.

N. Hjort and D. Pollard. Asymptotics for minimisers of convex processes. Unpublished manuscript, 1993. URL `http://www.stat.yale.edu/~pollard/Papers/convex.pdf`.

A. Kumar, S. Vembu, A. K. Menon, and C. Elkan. Beam search algorithms for multilabel learning. *Machine Learning*, 92(1), 2013.

G. Madjarov, D. Kocev, D. Gjorgjevikj, and S. Džeroski. An extensive experimental comparison of methods for multi-label learning. *Pattern Recognition*, 45(9):3084–3104, 2012.

J. Mielniczuk and H. Szymanowski. Selection consistency of generalized information criterion for sparse logistic model. In *Stochastic Models, Statistics and Their Applications*, pages 111–119, 2015.

J. Mielniczuk and P. Teisseyre. What do we choose when we err? Model selection and testing for misspecified logistic regression revisited. Accepted for publication, 2015.

E. Montañes, R. Senge, J. Barranquero, J. R. Quevedo, J. J. del Coz, and E. Hüllermeier. Dependent binary relevance models for multi-label classification. *Pattern Recognition*, 47(3):1494 – 1508, 2014.

R. Nishii. Asymptotic properties of criteria for selection of variables in multiple regression. *Annals of Statistics*, 12(2):758–765, 1984.

D. Preigbon. Goodness of Link Tests for Generalized Linear Models. *Applied Statistics*, 29 (1):15–24, 1980.

G. Qian and C. Field. Law of iterated logarithm and consistent model selection criterion in logistic regression. *Statistics and Probability Letters*, 56:101–112, 2002.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2013. URL `http://www.R-project.org/`.

J. Read, B. Pfahringer, G. Holles, and E. Frank. Classifier chains for multi-label classification. In *ECML/PKDD*, pages 254–269, 2009.

J. Read, Pfahringer B. , G. Holmes, and E. Frank. Classifier chains for multi-label classification. *Machine Learning*, 85:333–359, 2011.

R. Senge, J. J. del Coz, and Hüllermeier E. . On the problem of error propagation in classifier chains for multi-label classification. In *Conference of the German Classification Society on Data Analysis, Machine Learning and Knowledge Discovery*, 2012.

R. Stukel. Generalized Logistic Models. *American Statistical Association*, 83(402):426–431, 1988.

P. Teisseyre. *On some methods of model selection for linear and logistic regression*. PhD thesis, Institute of Computer Science, Polish Academy of Sciences, 2013. URL `http://www.ipipan.eu/~teisseyrep/PUBLICATIONS/Teisseyre_PhD.pdf`.

K. Trohidis, G. Tsoumakas, G. Kalliris, and I. Vlahavas. Multilabel classification of music into emotions. In *International Conference on Music Information Retrieval*, pages 325–330, 2008.

G. Tsoumakas and I. Katakis. Multilabel classification: an overview. *International Journal of Data Warehouse and Mining*, 3:1–13, 2007.

T. Zhang. Some sharp performance bounds for least squares regression with l_1 regularization. *The Annals of Statistics*, 77(5A):2109–2144, 2009.

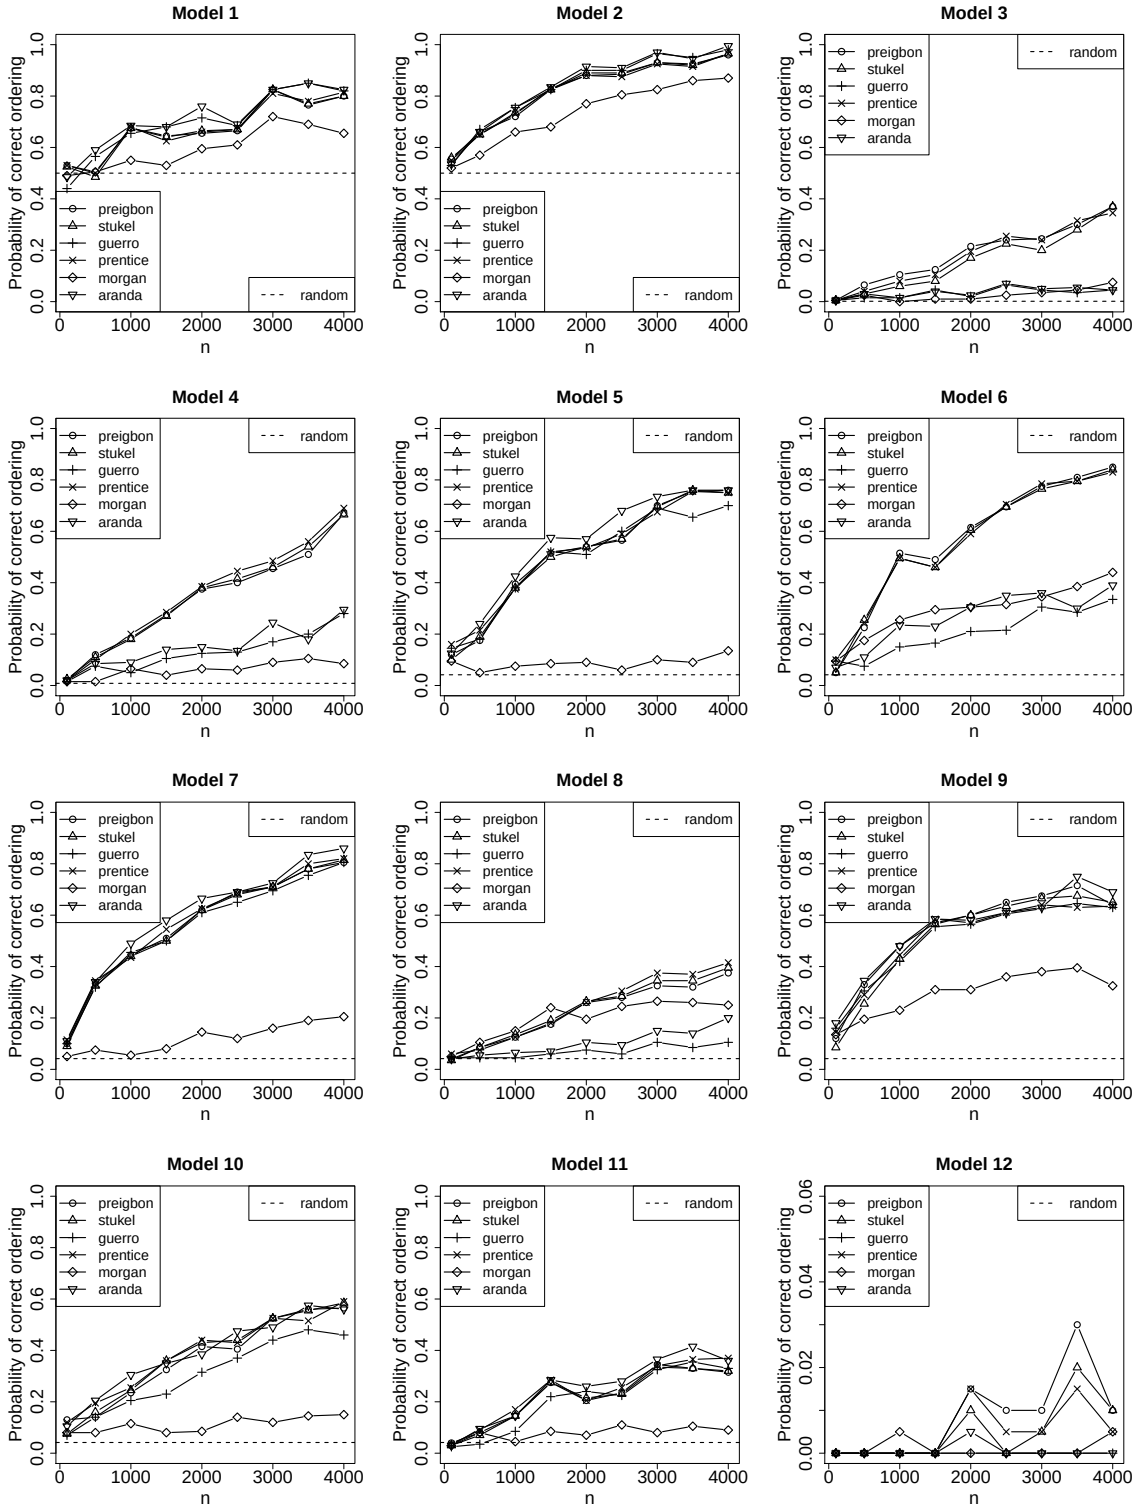

Fig. 4 Probabilities of finding the correct ordering with respect to the number of training cases. The scale for model (M12) is much smaller than for other models.

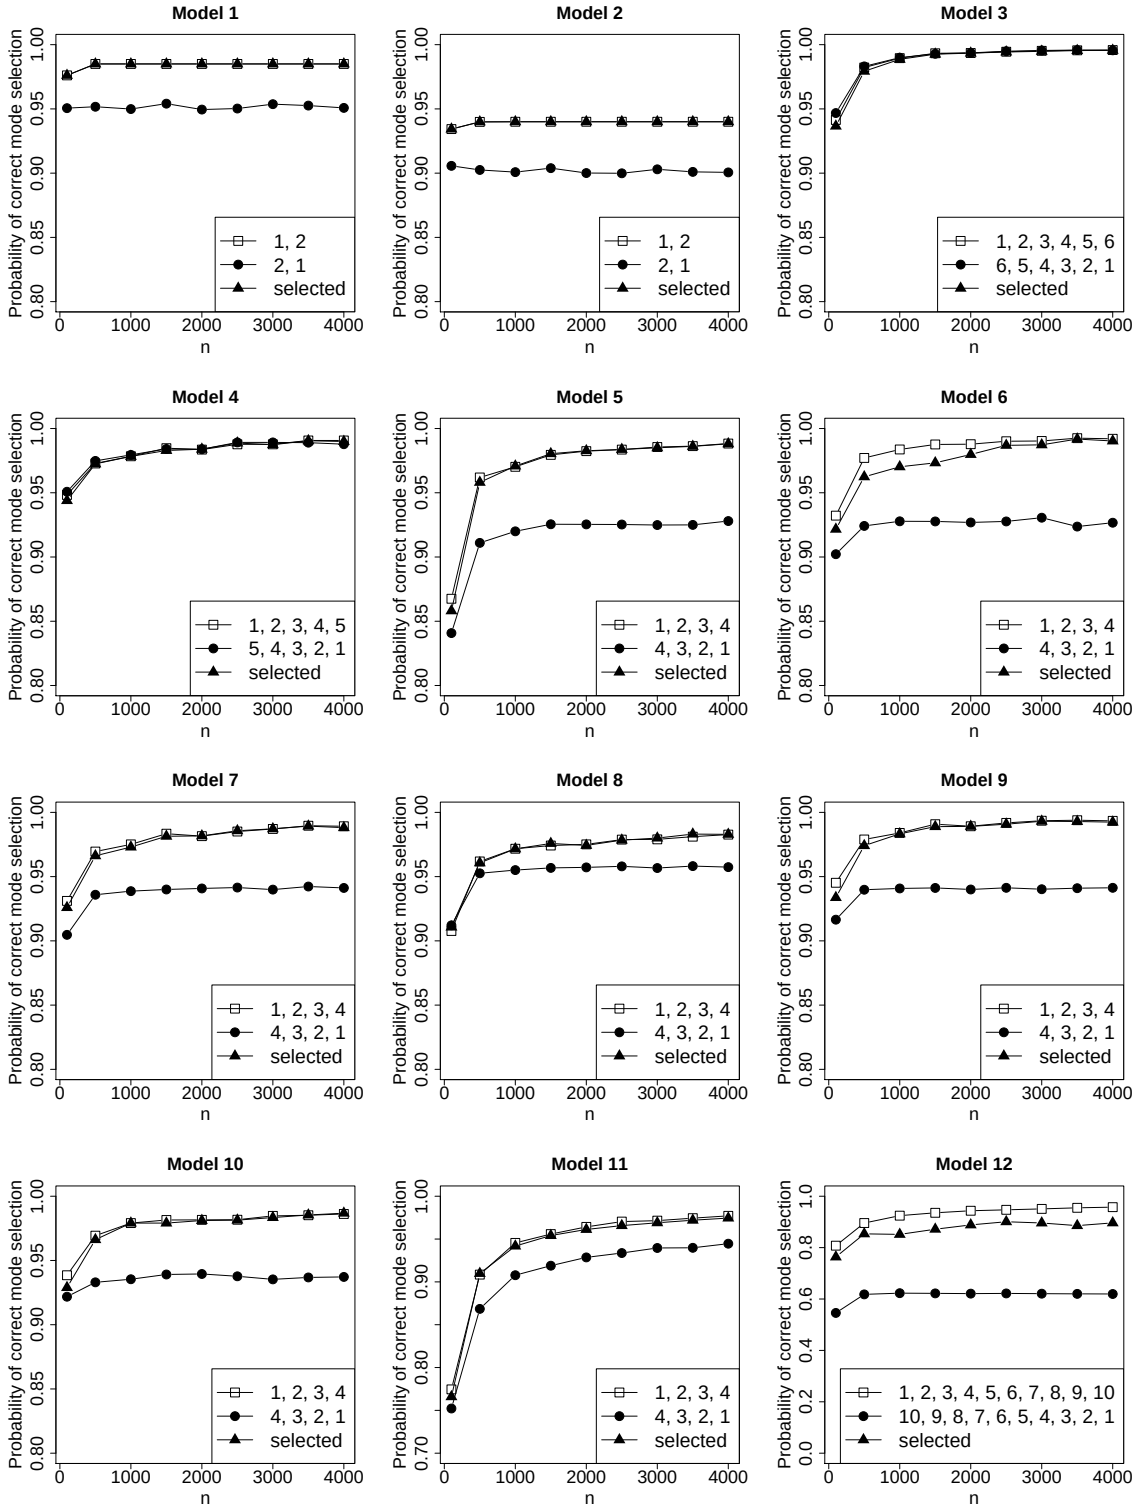

Fig. 5 Probabilities of correct mode selection with respect to the number of training cases.